Error reduction by parallel repetition – the state of the art

(Preliminary version)

Uriel Feige*

November 5, 1995

Abstract

We show that no fixed number of parallel repetitions suffices in order to reduce the error in two-prover one-round proof systems from one constant to another. Our results imply that the recent bounds proven by Ran Raz, showing that the number of rounds that suffice is inversely proportional to the answer length, are nearly best possible. Our proof technique builds upon an idea of Oleg Verbitsky.

We use this opportunity to survey the known results on parallel repetition, and to present the proofs of some previously claimed theorems.

1 Introduction

A two prover one round proof system [8], MIP(2,1), is a protocol by which two provers jointly try to convince a computationally limited probabilistic verifier that a common input belongs to a prespecified language. The verifier selects a pair of questions at random. Each prover sees only one of the two questions, and sends back an answer. The verifier evaluates a predicate on the common input and the two questions and answers, and accepts or rejects according to the output of the predicate. For inputs in the language, the provers have a strategy (where a strategy for a prover is a function from incoming messages to outgoing messages) that always causes the verifier to accept. For inputs not in the language, regardless of the strategy used by the provers, the verifier accepts with probability at most ϵ. The smaller the value of ϵ, known as the *error*, the more confidence the verifier has in using the proof system.

*Department of Applied Math and Computer Science, The Weizmann Institute, Rehovot 76100, Israel. feige@wisdom.weizmann.ac.il. Incumbent of the Joseph and Celia Reskin Career Development Chair. Supported by a Yigal Alon fellowship.

Reducing the error in MIP(2,1) proof systems is a subtle issue. A natural approach is to repeat an MIP(2,1) protocol n times, and accept only if all executions were accepting. Ideally, one would hope that this method reduces the error to ϵ^n. This is indeed true if each execution is performed with a fresh pair of provers, requiring n pairs of provers, or if the executions are performed sequentially (each prover must answer each question online, before seeing the question for the next execution), requiring n rounds of communication. However, parallel repetition, in which there are only two provers, and each prover sends out its answers only after receiving all its questions, is not guaranteed to reduce the error to ϵ^n [18]. Much work was invested in trying to analyse the rate at which parallel repetition reduces the error in MIP(2,1) proof systems.

There are many known results about error reduction by parallel repetition. The most useful of these is a theorem of Raz [26], that for any MIP(2,1) proof system with error ϵ there exists some constant α (that depends on the proof system) such that the error in n parallel repetitions is at most $\epsilon^{\alpha n}$. The main new result that we present is a proof that there is no universal α that fits all MIP(2,1) proof systems. Hence qualitatively, Raz's theorem tells the true story. (Quantitatively, there is still room for improvement, in giving the best α for each MIP(2,1) proof system.)

We feel that research on parallel repetition has reached a stage where the initial major open questions received satisfactory answers, and most questions that remain open are of a less fundamental nature. Since the proof of our main theorem is short, we use this opportunity to present an extensive survey of known results on parallel repetition. In addition, in an appendix to this paper, we present proofs of some theorems that were previously announced without proof. These proofs are included because we do not anticipate that they will be published anywhere else, and because they give a general idea of how other results are proved.

In Section 2 we formally define the parallel repetition problem. In Section 3 we survey the known results about parallel repetition. In Section 4 we state our main result, give a high level overview of its proof, and explain its relation to earlier work of Verbitsky [29] on which the proof builds upon. In Section 5 we give a more detailed statement of our main theorem, and present its proof. In Section 6 we discuss cases in which parallel repetition may not be the best choice for reducing the error. In the appendix we present previously unpublished proofs for some of the theorems cited in the survey.

2 Definitions

Two-prover one-round proof systems are often modeled as a game between a verifier and two cooperating provers. It is not an adversarial game, as the strategy of the verifier is fixed in advance. Rather, it is a cooperative game with two players (provers) who coordinate a joint strategy that gives them the highest probability of winning. Modelling MIP(2,1) proof systems as games abstracts away the language recognition aspects of these proof systems, but preserves the concept of error in a proof system, and the issue of how error can be reduced. One may best think of a game as an instantiation of an MIP(2,1) proof system on a single input that is not in the input language.

A (two-prover one-round) game $G = G(X, Y, Q, \pi, A, B, V)$ is a finite object with a question set $Q \subset X \times Y$, probability measure π induced on Q, answer set $A \bigcup B$, and acceptance predicate V. The *support* $Q \subset X \times Y$ is the set of question pairs that have positive probability of being asked, where X is the set of questions to prover P_1, and Y is the set of questions to prover P_2. The answer set $A \bigcup B$ is the union of two answer sets, where A is the set of answers available for prover P_1, and B is the set of answers available for prover P_2. The verifier V selects at random a question pair $(x, y) \in Q$, according to probability distribution π. Question x is sent to P_1 who replies with an answer $P_1(x) \in A$. Question y is sent to P_2 who replies with an answer $P_2(y) \in B$. (We identify between the name of a prover and the strategy that it employs.) The verifier then evaluates a predicate $V(x, y, P_1(x), P_2(y))$, and accepts if the predicate is satisfied. The goal of the provers is to select a strategy (namely, two functions, one for each prover, specifying an answer to each possible question) that maximizes the probability that the verifier accepts. The probability that the verifier accepts under the optimal strategy of the provers is denoted by $\omega(G)$. If $\omega(G) = 1$ the provers are said to have a *perfect strategy* for G, and the game G is *trivial.* For nontrivial games, $\omega(G)$ is also called the *error* of the game. We shall only be interested in nontrivial games.

An n-fold parallel repetition of game G is a new game, denoted by G^n, played on n coordinates. The verifier treats each coordinate of G^n as an independent copy of the original game G, and accepts in G^n only if it would have accepted all the n copies of G. The support of G^n set is Q^n, the answer set is $A^n \bigcup B^n$. The verifier selects n question pairs $(x_i, y_i) \in Q$ independently, each according to the probability distribution π. We denote $\bar{x} = x_1 x_2 \ldots x_n$ and $\bar{y} = y_1 y_2 \ldots y_n$. (Hence $\bar{x} \in X^n$, and $\bar{y} \in Y^n$.) A strategy for the provers is $2n$ functions, $P_1^i : X^n \longrightarrow A$ and $P_2^i : Y^n \longrightarrow A$, where $1 \leq i \leq n$. The acceptance predicate is $V^n = \bigwedge_{i=1}^n V(x_i, y_i, P_1^i(\bar{x}), P_2^i(\bar{y}))$. That is, the verifier accepts if all n copies of the original game G are accepting.

Observe that even though the verifier treats each coordinate of G^n independently, the provers may not. In particular, the answer a prover gives in coordinate i may depend on the questions that the prover receives in other coordinates. For this reason, it is not true in general that $\omega(G^n) = (\omega(G))^n$.

The game model of MIP(2,1) can be extended to allow for the acceptance predicate V to be a randomized predicate (that is, to have an additional input that is assigned randomly by the verifier). Known applications of MIP(2,1) proof systems do not make use of this extra option. Upper bounds on $\omega(G^n)$ were proved in the restricted model that we use, and the question of whether they hold in the extended model is most often ignored. Lower bounds that are proved in the restricted model hold also in the extended model.

Another extension of the MIP(2,1) model is to allow the provers to use randomized strategies. This is irrelevant for our purpose of analysing the error probability (since there always is a deterministic optimal strategy), but is important in the context of constructing *zero knowledge* MIP(2,1) proof systems. We shall not deal with zero knowledge issues.

We single out some classes of games:

1. **Uniform:** π is uniform on Q. Many (but not all) MIP(2,1) proof systems give rise to uniform games. Some techniques for analysing $\omega(G^n)$ apply only to uniform games (see Section 3.3). Bounds on $\omega(G^n)$ obtained for uniform games can be translated into (weaker) bounds for nonuniform games. Consider a nonuniform game and let p denote the lowest nonzero probability assigned to a question pair. The game can be treated as a sum of two games, one uniform, played with probability $p|Q|$, and the other nonuniform. If the original game was nontrivial, then so is its uniform version. When the original game is repeated n times, approximately $p|Q|n$ copies of it are of the uniform game, and the bounds for uniform games apply.

2. **Free:** The game is uniform and with full support (namely, $Q = X \times Y$). (The term *free* was introduced in [10]. Subsequently, the term games of *no information* was used in [16] to describe a slightly more general class of games, in which the questions to the two provers are chosen independently. Results proven for free games are often true for games of no information as well.) Some known results on parallel repetition apply only to free games, or are simpler to prove in this case. Unfortunately, free games are rare exceptions in MIP(2,1) proof systems (with [21] being the exception). This is partly explained by the fact that it is easy to approximate $\omega(G)$ for free games within an arbitrarily small additive factor (see [5]).

3. **Compare:** For this class of games, the acceptance predicate of the verifier can be decomposed into two parts, the first involving only x and $P_1(x)$, and the second involving x, y, $P_1(x)$, $P_2(y)$. The second part is of a special form, where $P_2(y)$ is compared for exact equality with a segment of $P_1(x)$, where the choice of the segment depends on x and y. Compare games arise naturely in MIP(2,1) proof systems. Moreover, the compare structure has been useful in certain applications (hardness of approximating SETCOVER [24], hardness of approximating CLIQUE [7]).

4. **Unique:** Given x and y, the acceptance predicate V induces a (partial) permutation on $P_1(x)$ and $P_2(y)$. That is, for any answer $P_1(x)$ there is at most one answer $P_2(y)$ that causes V to accept, and vice versa. Parallel repetition for these games is easier to analyse, but they do not arise naturely in MIP(2,1) proof systems, because it is

easy to tell if a unique game is trivial (see [16], and also [13] regarding the hardness of approximating $\omega(G)$ for unique games).

3 A survey

In this section we survey most results known about parallel repetition. It is not the purpose of this survey to give proper credit to all those who deserve credit. Neither is it our purpose to present the results in their correct historical order. Rather, it is our intention that this short survey will present the state of the art regarding error reduction in MIP(2,1) proof systems, implicitly point to questions that remain open, and provide references that contain ideas that may help future progress. The sooner the survey becomes outdated, the better it served its purpose.

Issues such as how MIP(2,1) proof systems are constructed, their applications to proving hardness of approximation results, zero knowledge versions and applications to cryptography, are not discussed in this survey.

3.1 Negative results

MIP(2,1) proof systems were introduced by Ben-Or, Goldwasser, Kilian, and Wigderson [8]. The authors remarked that parallel repetition of these proof systems preserves zero knowledge properties, but the issue of its effect on the error was not touched upon. Following an incorrect claim by Fortnow, Rompel and Sipser [17] that $\omega(G^n) = (\omega(G))^n$ (which remained unchallenged by many who read the preliminary version of that paper, including myself), Fortnow constructed an example showing $\omega(G^2) > (\omega(G))^2$ [18]. Lapidot and Shamir [21] showed the same effect in a natural zero knowledge MIP(2,1) proof system for NP. Feige [12] presented an example of a nontrivial game in which $\omega(G^2) = \omega(G)$. All the above examples are for free games. A nonfree (but uniform) game with $\omega(G^2) = \omega(G)$ is presented in [16].

For games with $k > 2$ provers, Raz (private communication, see Appendix A) shows an example where $\omega(G^k) = \omega(G)$.

3.2 Analysis based on recursion

The most useful bound on $\omega(G^n)$ was proved by Raz [26]. A rough and incomplete sketch of the proof technique (that overlooks many of the subtleties involved) follows. For simplicity of this presentation, assume that G is a uniform game.

Fix the optimal strategy for the provers on G^n, and let S^n be the set of question pairs on which the verifier accepts. We need to bound $|S^n| = \omega(G^n)|Q^n|$. We upper bound $|S^n|$ as follows. Let set U be a tentative candidate for the set of question pairs on which the verifier accepts. Initially, let $U = Q^n$. If U is sufficiently large, then entropy arguments are used in order to show that at least one coordinate i is *good*, where the properties of good coordinates are discussed shortly. On this good coordinate, it can be shown that the provers succeed with probability not much higher than $\omega(G)$. Partition U into two sets, U_{accept}, the set of question pairs accepted by the verifier, and U_{fail}, the the set question pairs rejected by the verifier. The set U_{fail} is removed from U, and for U_{accept} we attempt to find coordinates other than i that cause rejection.

Partition U_{accept} into $|Q|$ sets, denoted by U_k, where $1 \leq k \leq |Q|$. Set U_k contains all question pairs of U_{accept} for which on the good coordinate i, the question pair (x, y) asked is the kth question pair. For each set U_k we ignore the question pair on the ith coordinate. Hence if U was viewed as a subset of question pairs for the game G^r, then U is viewed as a subset of question pairs for the game G^{r-1}. Any startegy played by the provers on U transforms to a legal strategy for the provers on U_k, because in U_k, both provers can complete the question pair (x, y) of the missing coordinate i, and answer as the original strategy on U would. For technical reasons that make the proof go through, each set U_k is further partitioned into $|A|$ sets, denoted by $U_{k,\ell}$, where $1 \leq \ell \leq |A|$. Set $U_{k,\ell} \subset U_k$ contains the question pairs for which the answer that prover P_1 gave on the good coordinate i is the ℓth answer, and the answer that P_2 gave is one that causes acceptance. For each of the sets $U_{k,\ell}$, if it is sufficiently large, a new good coordinate j is selected, and the argument is repeated recursively. The final bound obtained on $\omega(G^n)$ depends on m, the number of times this argument can be repeated.

To get a rough estimate on m, observe that if $|U|/|Q^r| \leq (\omega(G))^r$, the provers need not fail, and the argument cannot be applied. On average, $|U_{k,\ell}|/|U| \simeq \omega(G)/(|Q| \times |A|)$. However, note that $|U_{k,\ell}|$ is compared against $|Q^{r-1}|$, whereas $|U|$ is compared against $|Q^r|$. Hence, when selecting a good coordinate, the effective reduction in size is only by a factor of $\omega(G)/|A|$. Thus m roughly satisfies $(\omega(G)/|A|)^m \simeq (\omega(G))^{n-m}$. We obtain that $m \simeq \alpha n$, where $\alpha \simeq \log(1/\omega(G))/(\log|A| + \log(1/\omega(G)))$. As we see, at best, this approach can prove bounds of the form $\omega(G^n) \leq (\omega(G))^{\alpha n}$, for α as defined above. In particular, the bounds deteriorate as $|A|$ grows, and as $\omega(G)$ approaches 1.

Raz presents his results using somewhat different terms. He shows the existence of a function W that maps from the open interval $(0,1)$ to itself (in fact, to the open interval $(t,1)$, where t is some constant, $0 < t < 1$), such that $\omega(G^n) \leq (W(\omega(G)))^{n/\log|A|}$. Raz does not explicitly describe W, but notes that for any fixed value of $\omega(G)$, the value of $W(\omega(G))$ is just another fixed constant, and hence the effect of W can be replaced by some constant in the exponent. The value of this constant depends on $\omega(G)$, and careful examination shows that it behaves like $(1-\omega(G))^c/\log(1+1/\omega(G))$, for some constant c. Qualitatively, the behavior of Raz's bound is similar to α as described above, except for the case where $\omega(G)$ is very small.

To complete our overview of Raz's proof, we explain how good coordinates are defined.

Recall that on a good coordinate, we want the error to be not much higher than $\omega(G)$. Hence a natural requirement is:

Property 1: The projection of U on coordinate i gives a probability distribution on the question pairs Q that is similar to π.

This property indeed suffices for free games. But for games that are not free, this is not a sufficient property. Recall that for Q^n, the question pairs selected at different coordinates are independent. This is no longer true for sets U obtained in the recursive proof described above. For example, assume that the answer that a prover gives on coordinate i is the sum modulo $|A|$ of the questions on coordinates j and k. Then after we end the recursive step on coordinate i, and partition U according to A, we are in the situation where the question the the prover receives on coordinate j is uniquely determined by the questions it receives on coordinate k. The situation may become even worse. If the game is not free, then the question a prover receives in coordinate k is corollated with the question that the other prover receives on this coordinate. Hence the question that one prover receives on coordinate k becomes corollated with the question that the other prover receives on coordinate j. This *side information* may help the provers succeed on coordinate j. This leads to the following property for good coordinates:

Property 2: The side information available to each prover (namely, the questions a prover receives on coordinates other than i) conveys almost no information on the question that the other prover receives on coordinate i (beyond the direct information available to the prover through its own question on coordinate i and the description of π).

By now, it should not come as a surprise that a coordinate may have both properties 1 and 2, and still the provers succeed with probability much higher than $\omega(G)$. An explicit example, in which coordinate i is perfect with respect to both properties above, yet the provers have a perfect strategy on i, is presented in Appendix B.

The above discussion motivates the central property that Raz requires from good coordinates. One way of formalizing this property (somewhat different from Raz's formulation) is as follows. The *direct information* a prover has on coordinate i is the question it receives on this coordinate. The direct information provides some information about the question that the other prover receives on coordinate i (if G is not a free game).

Property 3: Let the *extended side information for coordinate* i be the questions a prover receives on coordinates other than i, together with the questions that the other prover receives on half of the coordinates, where these coordinates are chosen at random, excluding i. Then on average, the extended side information conveys almost no additional information on the question that the other prover receives on coordinate i.

With properties 1 and 3 defining good coordinates, it is possible to show that if U is large enough then a good coordinate i exists, and that on this coordinate, the provers succeed with probability that is not much higher than $\omega(G)$.

In this context, it is worth mentioning a qualitative version of Raz's theorem, for free games. This is an unpublished result of Feige and Szegedy (cited in [14]) that for free games, $\omega(G^3) < \omega(G)$ (contrasting with the negative examples of section 3.1). The techniques used in the proof of this result are related to those used by Raz (though much simpler). See Appendix C for details.

3.3 The forbidden subgraph approach

Another approach for obtaining bounds on $\omega(G^n)$ is that of analysing an extremal graph theory type of question, which abstracts away the answer set. The extremal graph theory approach is applicable only for uniform games. The game G^n is represented as a bipartite graph, where vertices on the left side are labeled by X^n, and vertices on the right side are labeled by Y^n. If $(\bar{x}, \bar{y}) \in Q^n$, then there is an edge $(\bar{x}, \bar{y})$. A strategy of the provers corresponds to a two coloring of the edges of the bipartite graph, such that red edges correspond to question pairs on which the verifier accepts, and blue edges correspond to question pairs on which the verifier rejects. This coloring must be consistent with the assumption that G is nontrivial, or more generally, with the value of $\omega(G)$. This translates to *forbidden subgraphs* restrictions on the coloring. A forbidden subgraph has more than $\omega(G)|Q|$ red edges, and its vertices satisfy the following projection requirement: for some index i, the projection of the lefthand side vertices on i gives the whole of X with no repetitions, and the projection of the righthand side vertices on i gives the whole of Y with no repetitions. The existence of such a subgraph implies strategy for G that succeeds with probability greater than $\omega(G)$, a contradiction. If G is uniform, then $\omega(G^n)$ is upper bounded by the maximum fraction of edges that can be colored red without producing a forbidden subgraph.

The above approach was used in [10, 21, 25, 1, 12] to analyse nontrivial free games. For the special case that $|X| = |Y| = 2$, Peleg [25] shows that $\omega(G^n) = O(2^{-(1-(\log 3/2))n}) \simeq O(2^{-0.21n})$, and that the approach cannot give anything better than $\omega(G^n) = 2^{-n/3}$. For free games with arbitrary support size, Alon [1] and Feige [12] prove similar bounds, showing that $\omega(G^n) = 2^{-O(n/(|X|\cdot|Y|\cdot\log|X|))}$. Upto the $O(\log|X|)$ factor in the exponent, this result matches the trivial lower bound. If in addition, $\omega(G) < p^2$, then Feige [12] proves that $\omega(G^n) < e^{-k}$, where $k = n\log(1/p)/2(|X| + \log(1/p) + 2)$. For constant p, this matches the lower bound implied by Section 5, up to an $O(\log|X|)$ factor in the exponent.

Much less is known about nonfree games. Verbitsky [28] uses Ramsey Theory [19] to show that as $n \to \infty$, $\omega(G^n) \to 0$, without giving a constructive bound. Verbitsky [29] gives a constructive upper bound of $\omega(G^n) = 2^{-O(n/(|Q|)^8)}$ for a family of nonfree games, in which Q induces a tree structure on $X \bigcup Y$. In the same paper, Verbitsky also shows that for nontrivial uniform games, the extremal graph theory question *exactly* characterizes the best bounds that can be proven on $\omega(G^n)$, if these bounds are only to depend on Q (and not the answer length, or any structural property of the acceptance predicate). This issue will be discussed again in Section 4.

3.4 Bounds using semidefinite programming

A different approach to bound $\omega(G^n)$ was investigated by Feige and Lovasz [16]. One introduces a function σ that has the following two properties:

1. For any game G, $\omega(G) \leq \sigma(G) \leq 1$.
2. For any game G and any n, $\sigma(G^n) = (\sigma(G))^n$.

From the above two properties it follows that $\omega(G^n) \leq (\sigma(G))^n$. Feige and Lovasz introduced such a function $\sigma(G)$ whose value is a solution to a semidefinite relaxation of a quadratic programming formulation of $\omega(G)$ (this particular function was inspired by [23]). For some classes of games, $\omega(G) < 1$ iff $\sigma(G) < 1$, implying that parallel repetition reduces the error at an exponential rate. In particular, this was proved for unique games [16], and for any game in which $|A| = |B| = 2$ [13]. Unfortunately, in other cases, it may happen that $\omega(G) < \sigma(G) = 1$, making the $\sigma(G)$ bound useless.

3.5 Making use of special structure

For games that have special structure, it is possible to give bounds that are independent of question and answer size. One such example (claimed without proof in [16]) is that of unique free games, for which $\omega(G^n) \leq (\omega(G))^{n/4}$. A proof of this appears in Appendix D.

Feige and Kilian [14] (extending a line of research studied in [9, 20, 12]) considered a class of games that are called "confuse or compare" games. In these games, with probability 1/2, the verifier plays a compare game (in which the number of segments involved is a constant), and with probability 1/2, the verifier plays a free game. They show that for confuse or compare games, the number of parallel repetitions needed in order to reduce the error from p to q is at most $(1/q(1-p))^c$, for some universal constant c.

3.6 Algebraic error reduction techniques

Algebraic error reduction techniques either use an algebraic structure that certain games derived from MIP(2,1) proof systems have (as in [11, 22]), or transform a game G into a new game $\hat{G}$ in which such a structure exists (as in [16]). Recall that the difficulty of analysing parallel repetition stems from the fact that a prover can base its answer to a question on questions in other coordinates of the game. The algebraic structure can be used in order to overcome this difficulty, by essentially forcing a prover to treat each of the coordinates independently. This typically involves the following mechanism. A prover is viewed as a function from questions to answers. Both the domain and the range of the function are extended in a certain *multilinear* way. In each coordinate, the first prover receives a question q and "line" ℓ in the extended domain that goes through q and has a random direction. The first prover sends a low degree polynomial over some finite field F that specifies the values of the extended function on all points of the ℓ, and in particular on q. The second prover is requested to send the value of the extended function on a random point y, where y lies on ℓ. The verifier checks (among other things) that the answer of the second prover is consistent with the polynomial of the first prover. The analysis of such proof systems uses the property that the point y is not correlated with the question q (due to the random direction of line ℓ), and the fact that two different low degree polynomials must differ on almost all their points. If the first prover is inconsistent in the answers it gives to q (basing the answer on information from other coordinates), this forces different polynomials to be sent, and inconsistencies arise on almost all possible values of y. Using this special structure of such games, it is relatively easy (compared to games that do not have such structure) to prove exponentially small bounds on $\omega(G^n)$, up to values that are inversely proportional to $|F|$.

4 The rate of error reduction depends on the answer size

Raz has shown [26] that for any game G, parallel repetition reduces the error at exponential rate. More specifically, $\omega(G^n) \leq (\omega(G))^{\alpha n}$, for some constant α that depends on G. Explicitly,

$$\alpha = \Omega\left(\frac{(1-\omega(G))^c}{\log |C_A| \cdot \log(1 + 1/\omega(G))}\right)$$

where c is some universal constant, and C_A is a parameter that involves covering the answer set A by generalized rectangles, and is never more than $|A|$. The most disturbing feature about this bound is that α is not a fixed constant (unlike the case of sequential repetition, where the constant in the exponent is 1). We show that Raz's bound is qualitatively optimal, in the sense that a dependency of α on G is unavoidable.

Theorem 1 *For any constant $\alpha > 0$, there are games G, such that for large enough n, $\omega(G^n) > (\omega(G))^{\alpha n}$.*

A detailed description of the properties of the games G above appears in the beginning of Section 5.

We do not know whether Raz's bound is quantitatively optimal. We only give examples in which $\alpha = \Omega(\log\log |A|/\log |A|)$, without resolving the question of whether α should depend on $\omega(G)$.

The proof of our theorem was inspired by a theorem of Verbitsky [29] that showed that the forbidden subgraph approach (described in section 3.3) is best possible if the only available information about the game G is a description of Q (the question-pair set), and that it is uniform and nontrivial. Verbitsky's theorem implies that negative results in the forbidden subgraph framework (such as the $2^{-n/3}$ lower bound in [25]) apply also to the parallel repetition problem. This would lead to a direct proof of Theorem 1, if we knew of strong enough lower bounds for the forbidden subgraph problem. To make the task of proving strong enough lower bounds easier (or even possible), we prove a theorem similar to that of Verbitsky, but for games in which $\omega(G) \simeq 1/2$. Unlike the case of nontrivial games, we do not get an exact correspondance between error reduction by parallel repetition and the forbidden subgraph problem. But the approximate correspondence that we get, together with a lower bound on the forbidden subgraphs problem, suffice.

We give a high level overview of the proof of Theorem 1. For any n, we exhibit a game for which n repetitions are required in order to reduce the error from one constant to another. This implies our theorem. First, we partially define a free game G, specifying only the size of $|Q|$, without defining the answer set or the acceptance predicate. Then we turn to the forbidden subgraph problem for G^n, in which we wish to color red as many edges as possible without having a forbidden subgraph with $|Q|/2$ red edges. A probabilistic argument shows that a constant fraction of the edges can be colored red (for an appropriate choice of $|Q|$ as a function of n). Now we return to the game G and define its answer set and acceptance predicate. This is done in a way that is consistent with the pattern of edges that are colored red in the forbidden subgraph problem. Finally, we prove that for G defined in this way, $\omega(G) \leq 3/4$.

5 The lower bound

In this section we show the existence of an infinite family of games $\{G_k\}$ with the following properties.

1. The support of game G_k is $X_k \times Y_k$, where $X_k = \{x_1, \ldots, x_k\}$ and $Y_k = \{y_1, \ldots y_k\}$. Hence the size of the support is k^2. For simplicitly of notation, we shall omit the subscript k from X_k and Y_k, as the value of k will be clear from the context.

2. Each game G_k is free. That is, the verifier picks a question pair uniformly at random from $X \times Y$.

3. The answer length of game G_k is $\Theta(k)$. (Hence length of answer is exponential in length of question.)

4. The error of game G_k is constant. More specifically, $\omega(G_k) \leq 3/4$.

5. For all $n \le k/4\log k$, the error in the n-fold parallel repetition of G_k remains above some constant. More specifically, $\omega((G_k)^n) \ge 1/8$.

For arbitrary large enough k, we explain how to obtain G_k with the properties described above. The definition of the acceptance criteria in game G_k depends on a certain graph G_k^n, that will be described shortly. Recall that in G_k, the question set for prover P_1 is $X = \{x_1, \ldots, x_k\}$, and the question set for prover P_2 is $Y = \{y_1, \ldots, y_k\}$. To describe the answer set of G_k, let $n = k/4\log k$ (rounded down to the nearest integer). The answer set for prover P_1 is $\{1, \ldots, n\} \circ X^n$ (that is, an integer from 1 to n, followed by a list of n questions x_i), and the answer set for prover P_2 is $\{1, \ldots, n\} \circ Y^n$. The verifier picks a question pair (x', y') uniformly at random from $X \times Y$. The verifier accepts the answer pair $(\ell \circ x_{i_1} x_{i_2} \ldots x_{i_n})$, $(m \circ y_{j_1} y_{j_2} \ldots y_{j_n})$, if and only if $x' = x_{i_\ell}$, $y' = y_{i_m}$, $\ell = m$, and the two vertices named $x_{i_1} \ldots x_{i_n}$ and $y_{j_1} \ldots y_{j_n}$ are connected by an edge in the bipartite graph G_k^n, to be described below.

Each side of the bipartite graph G_k^n has k^n vertices, where lefthand side vertices are labeled by the strings X^n, and righthand side vertices are labeled by the strings Y^n. The bipartite graph has $k^{2n}/8$ edges arranged in such a way that any k by k vertex induced subgraph of G_k^n has at most $k^2/2$ edges.

Lemma 2 *A graph G_k^n with the parameters above exists.*

Proof: We show that a random graph selected from an appropriate distribution has positive probability of satisfying the requirements of G_k^n (an existence proof by a probabilistic argument). The lefthand side vertex set is X^n and the righthand side vertex set is Y^n. Include each edge (connecting different sides) in the graph independently with probability 1/8. Then the expected total number of edges is $k^{2n}/8$, and with probability roughly 1/2 the number of edges will be at least its expectation. Consider now an arbitrary k by k vertex induced subgraph. The probability that it has $k^2/2$ edges (or more) is less than $8^{-k^2/2}\binom{k^2}{k^2/2} \le 2^{-k^2/2}/4$ (for large enough k). There are less than k^{2kn} ways of choosing a k by k vertex induced subgraph. Hence, if $n \le k/4\log k$, the probability that at least one of them has $k^2/2$ edges is at most 1/4. $\square$

To make the description of G_k concrete, we fix G_k^n to be the first (under an arbitrary order) bipartite graph on k^n by k^n vertices that satisfies the restrictions specified for G_k^n.

Lemma 3 *For the game G_k and n as described above, $\omega((G_k)^n) \ge 1/8$.*

Proof: We describe a strategy of the provers for the game $(G_k)^n$. On receiving a question $x_{i_1} \ldots x_{i_n}$, P_1 gives the n answers $\ell \circ x_{i_1} \ldots x_{i_n}$, for $\ell = 1, \ldots, n$. On receiving a question $y_{j_1} \ldots y_{j_n}$, P_2 gives the n answers $m \circ y_{j_1} \ldots y_{j_n}$, for $m = 1, \ldots, n$. By the definition of G_k, the strategy of the provers succeeds on all n repetitions of G_k simultaneously, whenever $(x_{i_1} \ldots x_{i_n}, y_{j_1} \ldots y_{j_n})$ is an edge of G_k^n. $\square$

Lemma 4 *For the game G_k as described above, $\omega(G_k) \le 3/4$.*

Proof: Consider an arbitrary strategy for the two provers for G_k. For each question x_i to P_1, this fixes an answer $\ell_i \circ x_{i_1} \dots x_{i_n}$. For each question y_j to P_2, this fixes an answer $m_j \circ y_{j_1} \dots y_{j_n}$. Construct the bipartite graph G with lefthand side vertex set X, righthand side vertex set Y, and an edge between x_i and y_j if the verifier accepts $(x_i, y_j, P_1(x_i), P_2(y_j))$.

Proposition 1 *If x_i and y_j are in the same connected component of G, then $\ell_i = m_j$.*

Proof: Follow the path connecting x_i and y_j in G. Since each edge (x', y') along the path signifies the fact that V accepts the provers' answers, for both endpoints of the edge the answers of the provers on x' and y' must respect $\ell' = m'$. The proof follows by transitivity. □

Proposition 2 *Any connected component of G has at most $k^2/2$ edges.*

Proof: Each answer of a prover is composed from an index (denoted by ℓ or m) and a name of a vertex in G_k^n. Within a connected component, all answers of the provers correspond to pairwise distinct vertices in G_k^n. This follows from Proposition 1, together with the fact that in order for the verifier to accept P_1's answer $m \circ x_{i_1} \dots x_{i_n}$ on x', it must hold that $x_{i_m} = x'$ (a similar condition holds for P_2). Each connected component can be viewed as (part of) a k by k vertex induced subgraph of G_k^n, and hence has at most $k^2/2$ edges (by definition of G_k^n). □

In Proposition 3 we prefer simplicity of proof over obtaining the best constant (which is $2 - \sqrt{2}$ rather than $3/4$).

Proposition 3 *G has at most $3k^2/4$ edges.*

Proof: For connencted component C_i in G, let a_i denote the number of vertices on the lefthand side and let b_i denote the number of vertices on the righthand side. Arrange the connected components C_i of G in order of decreasing b_i. C_1 contributes at most $k^2/2$ edges to G (by Proposition 2). For each C_i, $i > 1$, we compare between the number of edges that it adds to G and the number of edges that it excludes from being in G. Component C_i contributes at most $a_i b_i$ edges to G. Now count the number of edges that are excluded from G, by the fact that they join a vertex of C_i with a vertex of C_{i-1}. This number is $a_i b_{i-1} + a_{i-1} b_i > a_i b_i$ (because $b_{i-1} \ge b_i$). Hence beyond the first connected component, for any edge that we count in G we can count another edge as not being in G. □

Hence any strategy of the provers succeeds on at most $3/4$ of the support of G_k, and this completes the proof of Lemma 4. □

6 Discussion

Following Raz's theorem, parallel repetition establishes itself as the method of choice for reducing the error in MIP(2,1) proof system. Not only does it provably reduce the error at an exponential rate, but it is also preserves useful properties of MIP(2,1) proof systems, such as the compare structure, zero knowledge properties, and the power of the provers. For specific MIP(2,1) proof systems, in order to get the best possible bounds, one should bear in mind that there is a variety of techniques for analysing parallel repetition, and one can consult our survey to select the technique best fitted for the proof system.

The issue of parallel repetition is interesting in its own right, regardless of potential applications in the context of MIP(2,1) proof systems. But since the applications are also interesting, regardless of the issue of parallel repetition, we discuss several alternatives to employing straight parallel repetition.

- Algebraic error reduction techniques (see Section 3.6) give additional structure to the protocol that makes the analysis of parallel repetition easier. In some aspects, this extra structure is a disadvantage (e.g., it causes computational overhead for the provers, and overhead in the communication). It other cases, it is an advantage. In [2] this extra structure was used in a reduction proving the hardness of approximating a problem on lattices. In [14] this extra structure was used in order to simplify the design of zero knowledge proof systems.

  Algebraic error reduction techniques are applicable also in cases where one wants to transform MIP(k,1) proof systems (with a large number of provers) to an MIP(2,1) proof system, while preserving the error [16]. In some cases, they were used in the construction of proof systems with more than two provers [6, 27], where the extra provers are used in order to reduce the complexity of other measures of the proof system.

- For some applications one would like an error reduction procedure in which the verifier uses as few random bits as possible. Parallel repetition has fundumental limitations in this respect, whereas some other approaches may not have the same limitations. See [15].

- In some cases, simple transformations on MIP(2,1) proof system prior to employing parallel repetition may result in a more efficient error reduction procedure. One such example is in [11], where using a finite field that is larger than that necessary for a single execution of the proof system results in tighter bounds on the error under parallel repetition. We describe two other methods of helping parallel repetition achieve fast error reduction.

One method is to change a game with partial support into a game of full support, by extending the predicate V to accept whenever the question pair (x, y) does not belong to the original support Q, regardless of the answers of the provers. Though this increases the error in a single execution of the protocol, it may reduce the error for multiple executions, as demonstrated in [16].

Another method is useful if one needs to reduce the error from one constant to another. In this case, as Theorem 1 proves, parallel repetition may be arbitrarily slow. However, any game can easily be changed into a confuse or compare game, and thereafter the error is reduced at a rate that is independent of any other parameter of the game (see [14] for details). This method uses only constant overhead in reducing the error from one constant to another in any MIP(2,1) proof system.

Though we have learned quite a lot about parallel repetition, we do not have a good characterization that for every game tells us the rate by which parallel repetition reduces its error. Indeed, certain aspects of the problem suggest inherent difficulties. The error $\omega(G)$ itself does not seem to have a good characterization, in the sense that it is NP-hard to compute it, given a description of G. Moreover, it is NP-hard even to distinguish between trivial games and games with $\omega(G) < \epsilon$, for arbitrarily small ϵ (proved in [14] in combination with [3], with earlier indications of a weaker version of this result already implicit in [4]). Hence approaches (as in Section 3.4) that try to bound $\omega(G^n)$ by first obtaining a good characterization of $\omega(G)$, or of some $\sigma(G)$ that satisfies $\omega(G) \leq \sigma(G) < 1$, need considerable sophistication in order to have any hope of success. Another property that spells difficulties is the high sensitivity to the quality of the random bits that the verifier uses: if the verifier selects the question pairs pseudo-randomly rather than truly at random, then parallel repetition does not reduce the error [15]. In contrast, in many other contexts for which repetition is known to reduce the error (the class BPP being among the simplest examples), pseudo-randomness can replace true randomness. Further evidence to the complexity of parallel repetition is provided by Theorem 1, that there is no universal constant α (independent of parameters of the game) that bounds the rate of error reduction.

Acknowledgements

I thank Ran Raz for many helpful discussions, and for his comments on an earlier version of this manuscript. The contents of the appendix are based on past work of myself and of others, and are presented with the permission of Laszlo Lovasz, Ran Raz, and Mario Szegedy.

References

[1] N. Alon, "Probabilistic Methods in Extremal Finite Set Theory", in: *Extremal Problems for Finite Sets, (P. Frankl, Z. Füredi, G. O. H. Katona and D. Miklós Eds.), Bolyai Society Mathematical Studies,3, Visegrád, Hungary, 1991, 39–57, 1991.*

[2] S. Arora, L. Babai, J. Stern, Z. Sweedyk, "The hardness of approximate optima in lattices, codes, and systems of linear equations", *Proc. of 34th Annual Symposium on Foundations of Computer Science, 724–733, 1993.*

[3] S. Arora, C. Lund, R. Motwani, M. Sudan, M. Szegedy, "Proof verification and hardness of approximation problems", *Proc. of 33rd Annual Symposium on Foundations of Computer Science, 14–23, 1992.*

[4] L. Babai, L. Fortnow, C. Lund, "Non-deterministic exponential time has two-prover intertactive protocols", *Computational Complexity, 1:3–40, 1991.*

[5] M. Bellare, U. Feige, J. Kilian, "On the role of shared randomness in two prover proof systems", *Proc. of 3rd Israel Symposium on the Theory of Computing and Systems, 199–208, 1995.*

[6] M. Bellare, S. Goldwasser, C. Lund, A. Russell, "Efficient probabilistic checkable proofs and applications to approximation", *Proc. of 25th Annual Symposium on the Theory of Computing, 294–304, 1993.*

[7] M. Bellare, O. Goldreich, M. Sudan, "Free bits, PCPs and non-approximability – towards tight results", in *Proc. of 36th Annual Symposium on Foundations of Computer Science, 422–431, 1995.*

[8] M. Ben-or, S. Goldwasser, J. Kilian, A. Wigderson, "Multi Prover Interactive Proofs: How to Remove Intractability", *Proc. of 20th Annual Symposium on the Theory of Computing, 113–131, 1988.*

[9] J. Cai, A. Condon, R. Lipton, "On Bounded Round Multi-Prover Interactive Proof Systems", *Proc. of 5th IEEE Symposium on Structure in Complexity Theory, 45–54, 1990.*

[10] J. Cai, A. Condon, R. Lipton, "Playing Games of Incomplete Information", *Theoretical Computer Science, 103, 25–38, 1992.*

[11] J. Cai, A. Condon, R. Lipton, "PSPACE is Provable by Two Provers in One Round", *Journal of Computer and System Sciences, 48, 183–193, 1994.*

[12] U. Feige, "On the Success Probability of the Two Provers in One Round Proof Systems", *Proc. of 6th IEEE Symposium on Structure in Complexity Theory, 116–123, 1991.*

[13] U. Feige, M. Goemans, "Approximating the value of two prover proof systems, with applications to Max 2SAT and Max DICUT", *Proc. of 3rd Israel Symposium on the Theory of Computing and Systems, 182–189, 1995.*

[14] U. Feige, J. Kilian. "Two prover protocols – low error at affordable rates", *Proc. of 26th Annual Symposium on the Theory of Computing, 172–183, 1994.*

[15] U. Feige, J. Kilian. "Impossibility results for recycling random bits in two prover proof systems", *Proc. of 27th Annual Symposium on the Theory of Computing, (1995) 457–468.*

[16] U. Feige, L. Lovasz, "Two-prover one-round proof systems, their power and their problems", *Proc. of 24th Annual Symposium on the Theory of Computing, 733–744, 1992.*

[17] L. Fortnow, J. Rompel, M. Sipser, "On the Power of Multi-Prover Interactive Protocols", *Proc. of 3rd IEEE Symposium on Structure in Complexity Theory, 156–161, 1988.*

[18] L. Fortnow, J. Rompel, M. Sipser, "Errata for On the Power of Multi-Prover Interactive Protocols", *Proc. of 5th IEEE Symposium on Structure in Complexity Theory, 318–319, 1990.*

[19] A. Furstenberg, Y. Katznelson, "A density version of the Hales-Jewett theorem", *Journal d'Analyse Mathematique 57, 64–119, 1991.*

[20] J. Kilian, "Strong Separation Models of Multi Prover Interactive Proofs", *DIMACS Workshop on Cryptography, October 1990.*

[21] D. Lapidot, A. Shamir, "A One-Round, Two-Prover, Zero-Knowledge Protocol for NP", *Combinatorica 15(2), 203–214, 1995.*

[22] D. Lapidot, A. Shamir, "Fully Parallelized Multi Prover Protocols for NEXP-time", *Proc. of 32nd Annual Symposium on Foundations of Computer Science, 13–18, 1991.*

[23] L. Lovasz, "On the Shannon Capacity of a graph", *IEEE Trans. on Information Theory, 25, 1–7, 1979.*

[24] C. Lund, M. Yannakakis, "On the hardness of approximating minimization problems", *JACM, 41(5), 960–981, 1994.*

[25] D. Peleg, "On the maximum density of 0–1 matrices with no forbidden rectangles", *Discrete Mathematics 140, 269–274, 1995.*

[26] R. Raz. "A parallel repetition theorem". *Proc. of 27th Annual Symposium on the Theory of Computing, 447–456, 1995.*

[27] G. Tardos. "Multi-prover encoding schemes, and 3-prover interactive proofs." *Proc. of 9th IEEE Symposium on Structure in Complexity Theory, 308–317, 1994.*

[28] O. Verbitsky, "Towards the Parallel Repetition Conjecture", *Proc. of 9th IEEE Symposium on Structure in Complexity Theory, 304–307, 1994.*

[29] O. Verbitsky, "The parallel repetition conjecture for trees is true", *Electronic Colloquim on Computational Complexity, TR95-013, 1995.*

APPENDIX

Certain theorems discussed in the survey never had their proofs published (or even written down). To make these proofs accessible to the reader, they are included (with permission of their authors) in this appendix.

A Multiple-prover parallel repetition

The following example was provided by Ran Raz (private communication). It generalizes the "noninteractive agreement" protocol of [12] to the case of k provers, for $k > 2$. Multiple-prover games are a natural extension of two prover games, and the following example can be understood without a formal definition of these proof systems.

Let m be an integer. The verifier selects k integers $z_1, \ldots, z_k$ independently at random from $\{0, \ldots, m-1\}$, and sends integer i to prover P_i. Each prover P_i has to reply with two numbers, $p_i \in \{1, \ldots, k\}$, and $v_i \in \{0, \ldots, m-1\}$.

The verifier accepts if all the following conditions hold:

1. For any two provers, $p_i = p_j$. (That is, all provers point to the same prover.)
2. Let $\ell = p_i$. (That is, all provers point to prover P_ℓ.) Then $v_\ell = z_\ell$.
3. $\sum_{i=1}^{k} v_i = 0 \pmod{m}$.

Claim 1 *For the above k-prover game, $\omega(G) = 1/m$.*

Proof: Fix an arbitrary strategy for the provers, and partition the executions of the game into $k+1$ disjoint classes. Class C_i, $1 \leq i \leq k$, contains all executions in which all provers, except possibly P_i itself, point to prover P_i. Class C_0 contains all other executions. By condition 1, the verifier rejects all executions of class C_0. As for any other class C_i, observe that in order for the verifier to accept, $v_i = z_i$. Fix the questions to all provers except for z_i. There is now only one out of m possible values for z_i on which condition 3 is satisfied. Hence the verifier accepts only on a $1/m$ fraction of C_i. $\square$

Claim 2 *For the above k-prover game, $\omega(G^k) = 1/m$.*

Proof: Upon receipt of k numbers $z_i^1, \ldots, z_i^k$, prover P_i replies with pairs (p_i^j, v_i^j) for $1 \leq j \leq k$, where $p_i^j = j$, and $v_i^j = z_i^i$. The provers succeed simultaneously on all rounds if $\sum_{i=1}^{k} z_i^i = 0 \pmod{m}$, which happens with probability $1/m$. $\square$

B Properties of good coordinates

We construct an example in which the provers have a perfect strategy on coordinates that satisfy properties 1 and 2 of Section 3.2. To convey the intuition behind the example, we construct it in stages.

We start with a free game G_1 in which $X_1 = Y_1 = A_1 = B_1 = \{0,1\}$. The acceptance predicate on question pair (x, y) and on answers a and b is $x \wedge y = a \oplus b$ (where $\wedge$ denotes *logical and*, and $\oplus$ denotes *exclusive or*). For game $(G_1)^2$ consider $S \subset (Q_1)^2$ that includes the following eight question pairs, written as (x_1x_2, y_1y_2).

Set S: $(00, 00)$, $(01, 01)$, $(00, 10)$, $(01, 11)$, $(10, 00)$, $(11, 01)$, $(10, 11)$, $(11, 10)$.

Proposition 4 *The projection of set S on its first coordinate gives the uniform distribution over question pairs (x, y). Moreover, the question y_1 is independent of the two questions x_1x_2, and the question x_1 is independent of the two questions y_1y_2. (Hence properties 1 and 2 hold.) Nevertheless, the provers have a strategy that never fails on the first coordinate.*

Proof: The fact that properties 1 and 2 hold can be verified by inspection. To describe the perfect strategy of the provers, observe that for set S we have that $x_1 \wedge y_1 = x_2 \oplus y_2$. Hence on questions x_1x_2, the answer of the first prover on the first coordinate is x_2. On questions y_1y_2, the answer of the second prover on the first coordinate is y_2. $\square$

Despite Proposition 4, game G_1 cannot serve as a complete example. The set S cannot be generated on two coordinates of $(G_1)^n$ as the result of the recursive procedure described in Section 3.2. In fact, for game G_1, which is a free game, property 1 suffices in order to define good coordinates.

We extend G_1 to a game G of partial support. The game G may be viewed as the disjoint union of two games: G_1 (described above), and G_2. With probability 1/2, the verifier chooses a question pair from G_1, and with probability 1/2, a question pair from G_2. For game G_2, the questions to each prover are $X_2 = Y_2 = \{\bar{0}0, \bar{0}1, \bar{1}0, \bar{1}1\}$. The game is uniform over its support Q_2, where the support is exactly described by set S above, in which a pair of bits $z_1 z_2$ is interpreted as $z_1\bar{}z_2$. Hence $\pi(G)$ assigns probability 2/16 for each of the four question pairs of G_1, and probability 1/16 to each of the eight question pairs of G_2. For game G_2, the verifier accepts regardless of the provers' answers.

For game G^3, define the following two sets of triples of questions.

Set S_1**:** $0\,0\,\bar{0}0$, $0\,1\,\bar{0}1$, $1\,0\,\bar{1}0$, $1\,1\,\bar{1}1$.

Set S_2**:** $\bar{0}0\,0\,0$, $\bar{0}1\,0\,1$, $\bar{1}0\,1\,0$, $\bar{1}1\,1\,1$.

Based on the sets S_1 and S_2 we describe $U \subset Q^3$. Set U contains all question pairs that can be generated by each prover receiving a triple from $S_1 \bigcup S_2$. Out of the 64 possible combinations, only 16 are legal, due to the fact that G has only partial support. Eight legal question pairs can be constructed if each prover receives a triple from S_1. These pairs are "picked out" by the third coordinate and the partial support of G_2. Observe that for these eight pairs, on the first two coordinates we get exactly the set S. Eight other legal question pairs can be constructed if each prover receives a triple from S_2. These pairs are picked out by the first coordinate.

Proposition 5 *For set U, the first coordinate is distributed exactly according to π. Questions $x_2 x_3$ convey no information about question y_1 (beyond the information available by the value of x_1 and the fact that G has partial support). Questions $y_2 y_3$ convey no information about question x_1. The provers have a perfect strategy on the first coordinate of U.*

Unlike set S, set U can be generated in the process of applying the recursive proof on G^n. Assume that on coordinate 4, each prover gives answer 0 to question 0 if and only if the first three coordinates contain one of the eight combinations in sets S_1 and S_2. Assume that coordinate 4 is chosen as the first good coordinate. Then the corresponding set U_{accept} is partition according to Q and A. The set that corresponds to $x = y = a = b = 0$ has the property that on any sequence of question pairs on coordinates 5 to n, the first three coordinates contain eight copies of the set U. (The reason that we get eight copies of U and not just one is that out of the first four coordinates, three are of the game G_1, and each question pair in G_1 has probability 2/16 rather than 1/16 of being asked.)

C Three repetitions suffice for free games

The following theorem was proved by Feige and Szegedy (May, 1993). (The claim in [14] that Feige and Szegedy proved this theorem for any game with full support is apparently incorrect.)

Theorem 5 *For any nontrivial free game, $\omega(G^3) < \omega(G)$.*

We present a sequence of lemmas that are used in the proof of the theorem. These lemmas do not require G to be a free game.

Lemma 6 *For any game G, any n, and any strategy for G^n on which $\omega(G^n) = \omega(G)$, the following properties must hold:*

1. *Acceptance on the n copies of G is fully corollated. Whenever V accepts on one copy, it accepts on all copies.*
2. *For any $(x, y) \in Q$, the (weighted) fraction of question pairs with first coordinate equal to (x, y) that are accepted is exactly $\omega(G)$.*

Proof: To prove the first property, consider a strategy for G^n that seeks to maximize the expected number of copies on which it succeeds, rather than the probability of accepting all copies simultaneously. By linearity of expectation, the expected number of copies on which it succeeds is at most $\omega(G)n$. Hence if for some $(\bar{x}, \bar{y}) \in Q^n$ a strategy succeeds on some coordinates but not on all, then the fraction of question pairs in Q^n (weighted according to their respective probabilities) for which the strategy causes acceptance on all copies is strictly less than $\omega(G)$.

To prove the second property, observe that for the game G^n, if both provers are given the additional information of what the question pair (x, y) is on the first coordinate, then we remain with a game G^{n-1}. Clearly, $\omega(G^{n-1}) \leq \omega(G)$. If for some $(x, y) \in Q$, the (weighted) fraction of question pairs with first coordinate equal to (x, y) that are accepted is stricly smaller than $\omega(G)$, this cannot be compensated for by question pairs with a different (x', y') on their first coordinate. □

Lemma 7 *Let G be an arbitrary game with full support ($Q = X \times Y$), probability measure π, and error $\omega(G)$. Let $\bar{X} \subset X^n$, $\bar{Y} \subset Y^n$, and $\bar{Q} = \bar{X} \times \bar{Y}$. Assume that with some strategy for G^n, the provers never fail on any $(\bar{x}, \bar{y}) \in \bar{Q}$. Then the measure of $\bar{Q}$ induced by π^n (the probability that the question pair that the verifier selects in G^n happens to belong to $\bar{Q}$) is at most $(\omega(G))^n$.*

Proof: Fix a strategy of the provers that never fails on $\bar{Q}$. Modify it as follows. Let $x \in X$ be a question that appears on the first coordinate of some $\bar{x} \in \bar{X}$, and let a be the answer that the strategy gives on the first coordinate on question $\bar{x}$. Modify the strategy so that whenever x is the question on the first coordinate, then a is the answer on the first coordinate. The modified strategy never fails on $\bar{Q}$, because the answer a on question x never did (here we use the assumption that G has full support). We can continue in this fashion to fix particular answers to any question on the first coordinate, do the same for the second prover, and then repeat a similar process for every other coordinate. The final outcome of this process is a strategy that on each coordinate gives answers that are independent of the questions on the other coordinates, and never fails on $\bar{Q}$. But since strategies that treat each coordinate independently succeed with probability at most $(\omega(G))^n$, this limits the measure of $\bar{Q}$. $\Box$

A *probabilistic strategy* for a prover is a probabilistic function from questions to answers. That is, for question z, there is a probability associated with each answer a.

Lemma 8 *Let G be an arbitrary game and let P_1 and P_2 be a probabilitic strategy for the two provers that is optimal (succeeds with probability $\omega(G)$, where the probability of success is computed over the distribution of questions and the distribution of answers). Then for any question pair $(x, y) \in Q$, and for any pairwise distinct answers a_1 and a'_1 with $Pr[P_1(x) = a_1] > 0$ and $Pr[P_1(x) = a'_1] > 0$, and any pairwise distinct answers a_2 and a'_2 with $Pr[P_2(x) = a_2] > 0$ and $Pr[P_2(x) = a'_2] > 0$, the following* bipartite clique *condition holds:*

If $V(x, y, a_1, a_2)$ and $V(x, y, a'_1, a'_2)$ both accept, then $V(x, y, a_1, a'_2)$ and $V(x, y, a'_1, a_2)$ accept as well.

Proof: Assume that $V(x, y, a_1, a_2)$ and $V(x, y, a'_1, a'_2)$ both accept, but that $V(x, y, a_1, a'_2)$ does not. We present a probabilistic strategy that succeeds with probability greater than $\omega(G)$, a contradiction. This strategy is identical to the original strategy except for the following two changes.

Change 1: On question x, prover P_1 never answers with a'_1, and the probability that was associated with this answer is shifted to the answer a_1.

Change 2: On question y, prover P_2 never answers with a'_2, and the probability that was associated with this answer is shifted to the answer a_2.

Doing only one of the two changes cannot change the probability of success at all. This follows from the assumption that the original probabilistic strategy was optimal. Hence such a change cannot increase the success probability. It cannot decrease it either, because then the reverse change (never answering with a_1) would increase it.

Compare now the effect of Change 2 after Change 1 is done, to its effect before Change 1 is done. Change 2 involves question y and answer a'_2. The effect of Change 1 on evaluations of the predicate V that involve question y and answer a'_2 is limited to changing accepting instances $V(x, y, a'_1, a'_2)$ to rejecting instances $V(x, y, a_1, a'_2)$. Now Change 2 has the additional gain of transorming these new rejecting instances $V(x, y, a_1, a'_2)$ to accepting instances $V(x, y, a_1, a_2)$, a gain that is not present if Change 1 was not done beforehand.

Other then this, there is no influence of the Change 1 on the effect of Change 2. Hence the success probability increases beyond $\omega(G)$. $\square$

We are now ready to prove Theorem 5.

Proof: Let G be an arbitrary nontrivial free game with support $Q = X \times Y$, and with error $\omega(G)$. Assume for the sake of contradiction that $\omega(G^3) = \omega(G)$. Fix an optimal strategy P_1 and P_2 for the two provers of G^3. This fixes the pairs $(\bar{x}, \bar{y}) \in X^3 \times Y^3$ on which the verifier accepts.

Consider only the first coordinate of G^3, and the distribution of answers on it as induced by the strategies P_1 and P_2 for G^3. For each question $x \in X$ (and similarly, for questions $y \in Y$), this induces a probability distribution on the answers (where the probability is computed as the fraction of $x_1 x_2 \in X^2$ on which $P_1(x x_1 x_2)$ gives this answer on its first coordinate). This is a probabilistic strategy for G, since prover P_1 (and similarly P_2) can answer question x by first generating random questions x_1 and x_2, and then using the answer implied by $P_1(x x_1 x_2)$. Because G is a free game, there is no corollation between the randomness the two provers use in choosing their respective answers.

Consider any arbitrary question pair $q = (x, y)$ on the first coordinate. By Lemma 6, the probability of success on it is exactly $\omega(G)$. By Lemma 8, we have that the probabilistic strategy obeys the bipartite clique condition. Let a and b be the two sides of the bipartite clique. Consider now the game G^2 that remains on the other two coordinates by fixing q on the first coordinate. For this game, consider the set $\bar{X} \subset X^2$ that is induced by a and the set $\bar{Y} \subset Y^2$ that is induced by b, and let $\bar{Q} = \bar{X} \times \bar{Y}$. By Lemma 6, the provers strategy (conditoned on q being in the first coordinate) never fails on $\bar{Q}$. But now we have contradicting information regarding the measure of $\bar{Q}$. By Lemma 6, its measure is $\omega(G)$, whereas by Lemma 7, its measure can be at most $(\omega(G))^2$. $\square$

Remark: Though we have seen that for free games, $\omega(G^3) < \omega(G)$, Section 5 implies that $\omega(G^3)$ may be arbitrarily close to $\omega(G)$.

D Free games with the uniqueness property

The following theorem was announced by Feige and Lovasz [16].

Theorem 9 *For any unique game G that is free, $\omega(G^n) \leq (\omega(G))^{n/4}$.*

Before proving this theorem, we present a definition of games that are unique only with respect to one prover. In these games, once P_1 answers x, P_2 is forced to answer y in a unique way, or else the verifier rejects.

Definition 1 *A game G is* one-way unique *if for any $(x, y) \in Q$, and for any value s $(= P_1(x))$, there is at most one possible value t $(= P_2(y))$ which satisfies $V(x, y, s, t)$.*

The proof of theorem 9 is based on comparing between n parallel repetitions of G, and n sequential repetitions. To carry out the analysis we consider also a hybrid case, in which one prover receives all questions in parallel, and the other receives all its questions sequentially, and has to answer them online.

In the following lemma, we compare between $\omega_p(G^k)$, where p signifies the situation in which prover P_2 receives all questions in parallel, and $\omega_s(G^k)$, where s signifies the situation in which prover P_2 receives all questions sequentially. The lemma is applicable regardless of the mode of operation of prover P_1, whether parallel or sequential, as long as the same mode of operation is used in both cases that are considered. The proof of the lemma was inspired by a proof of Peleg in [25].

Lemma 10 *For any free one-way unique game G and any integer k, $\omega_p(G^n) \leq 2\sqrt{\omega_s(G^n)}$.*

Proof: Fix the optimal strategies P_1 and P_2 for G^n_p. This defines the following 0/1 matrix M: Each row X_i denotes a sequence of n questions to prover P_1, each column Y_j denotes a sequence of n questions to prover P_2, and an entry $M_{i,j}$ is '1' if $V^n(X_i, Y_j, P_1(X_i), P_2(Y_j))$ holds. Let r (c, respectively) denote the number of rows (columns) in M, and let $|M_1|$ denote the number of '1' entries in M. Clearly $|M_1|/rc = \omega_p(G^n)$. We bound $|M_1|$ from above.

Procedure *column delete*: Select row m with maximum number of '1' entries. Consider all the columns in which row m has '1'. Delete all these columns from M.

Claim 3 *The number of '1' entries deleted by Procedure* column delete *is at most $\omega_s(G^n)rc$.*

Proof: The answers of P_2 on each of the deleted columns succeeds with P_1's answers on row m. By one-way uniqueness, this means that P_2's answers in these columns are uniquely determined by P_1's respective answers on the maximum row. Hence prover P_2 can give the same sequence of answers also in a sequential mode of operation. $\Box$

Repeat Procedure *column delete* as long as there are rows left with more than $c\sqrt{\omega_s(G^n)}$ '1' entries. The number of repetitions is at most $1/\sqrt{\omega_s(G^n)}$, becuase afterwards there are no columns left. The total number of '1' entries deleted is thus at most $\omega_s(G^n)rc/\sqrt{\omega_s(G^n)}$. The total number of '1' entries which are left in the resulting matrix is at most $rc\sqrt{\omega_s(G^n)}$. It follows that M has at most $2rc\sqrt{\omega_s(G^n)}$ '1' entries. $\Box$

The proof of Lemma 10 extends to products of one-way unique games G in which π is a product distribution. We do not know if the lemma is true when π is arbitrary.

Games which are one-way unique with respect to prover P_1 are defined in an analogous way to Definition 1. If a game G is one-way unique with respect to both players, then G is has the *uniqueness* property (first introduced in [10]). Two successive applications of Lemma 10, first to "sequentialize" prover P_2 and then to "sequentialize" prover P_1, give the following corollary:

Corollary 11 *For any free game G with the uniqueness property and for any n, $\omega(G^n) \leq 2\sqrt{2}(\omega(G))^{n/4}$.*

To complete the proof of Theorem 9, we observe that the constant before $(\omega(G))^{n/4}$ (which is $2\sqrt{2}$ in the corollary) cannot be greater than 1. Assume for the sake of contradiction that for some G and n, $\omega(G^n) \geq (1+1/k)(\omega(G))^{n/4}$, for some positive integer k. Now consider $N = 2k$. By repeating the optimal strategy of G^n for $2k$ times, we obtain that

$$\omega(G^N) \geq ((1+1/k)(\omega(G))^{n/4})^{2k} \geq 4(\omega(G))^{N/4} > 2\sqrt{2}(\omega(G))^{N/4}$$

contradicting Corollary 11.

We do not know if Theorem 9 gives the best possible bound on free games with the uniqueness property. However, the bound in Lemma 10 is optimal (up to a factor of 2, but this factor can be shown not to exist if the mode of operation of prover P_1 is sequential). This can be deduced by examining the *noninteractive agreement* protocol, presented in Section A. For this protocol (with $k = m = 2$), $\omega(G^2) = \omega(G)$ even if the mode of operation of prover P_1 is sequential. We remark that it is known that for this noninteractive agreement protocol, $\omega(G^{2n}) = (\omega(G))^n$, because it is a game with the compare property holding simultaneously for both provers [12].